

An Improved Private Mechanism for Small Databases

Aleksandar Nikolov
 Microsoft Research
 Redmond, WA, USA
 alenik@microsoft.com

Abstract

We study the problem of answering a workload of linear queries \mathcal{Q} , on a database of size at most $n = o(|\mathcal{Q}|)$ drawn from a universe \mathcal{U} under the constraint of (approximate) differential privacy. Nikolov, Talwar, and Zhang [NTZ13] proposed an efficient mechanism that, for any given \mathcal{Q} and n , answers the queries with average error that is at most a factor polynomial in $\log |\mathcal{Q}|$ and $\log |\mathcal{U}|$ worse than the best possible. Here we improve on this guarantee and give a mechanism whose competitiveness ratio is at most polynomial in $\log n$ and $\log |\mathcal{U}|$, and has no dependence on $|\mathcal{Q}|$. Our mechanism is based on the projection mechanism of [NTZ13], but in place of an ad-hoc noise distribution, we use a distribution which is in a sense optimal for the projection mechanism, and analyze it using convex duality and the restricted invertibility principle.

1 Introduction

The central problem of private data analysis is to characterize to what extent it is possible to compute useful information from statistical data without compromising the privacy of the individuals represented in the dataset. In order to formulate this problem precisely, we need a database model and a definition of what it means to preserve privacy. Following prior work, we model a database as a multiset D of n elements from a universe \mathcal{U} , with each database element specifying the data of a single individual. Defining privacy is more subtle. A definition which has received considerable attention in recent years is *differential privacy*, which postulates that a randomized algorithm preserves privacy if its distribution on outputs is almost the same (in an appropriate metric) on any two input databases D and D' that differ in the data of at most a single individual. The formal definition is as follows:

Definition 1.1 ([DMNS06]). *Two databases D and D' are neighboring if the size of their symmetric difference is at most one. A randomized algorithm \mathcal{M} satisfies (ϵ, δ) -differential privacy if for any two neighboring databases D and D' and any measurable event S in the range of \mathcal{M} ,*

$$\mathbb{P}[\mathcal{M}(D) \in S] \leq e^\epsilon \mathbb{P}[\mathcal{M}(D') \in S] + \delta.$$

Differential privacy has a number of desirable properties: it is invariant under post-processing, the privacy loss degrades smoothly under (possibly adaptive) composition, and the privacy guarantees hold in the face of arbitrary side information. We will adopt it as our definition of choice in this paper. We will work in the regime $\delta > 0$, which is often called approximate differential privacy, to distinguish it from pure differential privacy, which is the case $\delta = 0$. Approximate differential privacy provides strong semantic guarantees when δ is $n^{-\omega(1)}$: roughly speaking, it implies that with probability at least $1 - O(n\sqrt{\delta})$, an arbitrarily informed adversary cannot guess from the output of the algorithm if any particular user is represented in the database. See [GKS08] for a precise formulation of this semantic guarantee.

We then turn to the question of understanding the constraints imposed by privacy on the kinds of computation we can perform. We focus on computing answers to a fundamental class of database queries: the *linear queries*, which generalize counting queries. A counting query counts the number of database elements that satisfy a given predicate; a linear query allows for weighted counts. Formally, a linear query is specified by a function $q: \mathcal{U} \rightarrow \mathbb{R}$ ($q: \mathcal{U} \rightarrow \{0, 1\}$ in the case of counting queries); slightly abusing notation,

we define the value of the query as $q(D) \triangleq \sum_{e \in D} q(e)$ (elements of D are counted with multiplicity). We call a set \mathcal{Q} of linear queries a *workload*, and an algorithm that answers a query workload a *mechanism*.

Since the work of Dinur and Nissim [DN03], it has been known that answering queries too accurately can lead to very dramatic privacy breaches, and this is true even for counting queries. For example, in [DN03, DMT07] it was shown that answering $\Omega(n)$ random counting queries with error per query $o(\sqrt{n})$ allows an adversary to reconstruct a very accurate representation of a database of size n , which contradicts any reasonable privacy notion. On the other hand, a simple mechanism that adds independent Gaussian noise to each query answer achieves (ϵ, δ) -differential privacy and answers any set \mathcal{Q} of counting queries with average error $O(\sqrt{|\mathcal{Q}|})$ [DN03, DN04, DMNS06].¹ While this is a useful guarantee for a small number of queries, it quickly loses value when $|\mathcal{Q}|$ is much larger than the database size, and becomes trivial for $\omega(n^2)$ queries. Nevertheless, since the seminal paper of Blum, Ligett and Roth [BLR08], a long line of work [DNR⁺09, DRV10, RR10, HR10, GHRU11, HLM12, GRU12] has shown that even when $|\mathcal{Q}| = \omega(n)$, more sophisticated private mechanisms can achieve error not much larger than $O(\sqrt{n})$. For instance, there exist (ϵ, δ) -differentially private mechanisms for linear queries that achieve average error $O(\sqrt{n} \log^{1/4} |\mathcal{U}|)$ [GRU12]. There are sets of counting queries for which this bound is tight up to factors polylogarithmic in the size of the database [BUV13].

Specific query workloads allow for error which is much better than the worst-case bounds. Some natural examples are queries counting the number of points in a line interval or a d -dimensional axis-aligned box [DNPR10, CSS10, XWG10], or a d -dimensional halfspace [MN12]. It is, therefore, desirable to have mechanisms whose error bounds adapt *both* to the query workload and to the database size. In particular, if $\text{opt}(n, \mathcal{Q})$ is the best possible average error² achievable under differential privacy for the workload \mathcal{Q} on databases of size at most n , we would like to have a mechanism with error at most a small factor larger than $\text{opt}(n, \mathcal{Q})$ for any n and \mathcal{Q} . The first result of this type is due to Nikolov, Talwar, and Zhang [NTZ13], who presented a mechanism running in time polynomial in $|\mathcal{U}|$, $|\mathcal{Q}|$, and n , with error at most $\text{polylog}(|\mathcal{Q}|, |\mathcal{U}|) \cdot \text{opt}(n, \mathcal{Q})$.

Here we improve the results from [NTZ13]:

Theorem 1.1 (Informal). *There exists a mechanism that, given a database of size n drawn from a universe \mathcal{U} , and a workload \mathcal{Q} of linear queries, runs in time polynomial in $|\mathcal{U}|$, $|\mathcal{Q}|$ and n , and has average error per query at most $\text{polylog}(n, |\mathcal{U}|) \cdot \text{opt}(n, \mathcal{Q})$.*

Notice that the competitiveness ratio in Theorem 1.1 is *independent of the number of queries*, which can be significantly larger than both n and $|\mathcal{U}|$. This type of guarantee is easier to prove when $n = \Omega(|\mathcal{Q}|)$, because in that case there exist nearly optimal mechanisms that are oblivious of the database size [NTZ13]. Therefore, we focus on the more challenging regime of small databases, i.e. $n = o(|\mathcal{Q}|)$.

It is worth making a couple of remarks about the strength of Theorem 1.1. First, in many applications the queries in \mathcal{Q} are represented compactly rather than by a truth table, and $|\mathcal{U}|$ is exponentially large in the size of a natural representation of the input. In such cases, running time bounds which are polynomial in $|\mathcal{U}|$ may be prohibitive. Nevertheless, our work still gives interesting information theoretic bounds on the optimal error, and, moreover, our mechanism can be a starting point for developing more efficient variants. Furthermore, under a plausible complexity theoretic hypothesis, our running time guarantee is the best one can hope for without making further assumptions on \mathcal{Q} [Ull13]. A second remark is that our optimal error guarantees are in terms of *average* error, while many papers in the literature consider worst-case error. Proving a result analogous to Theorem 1.1 for worst-case error remains an interesting open problem.

Techniques. Following the ideas of [NTZ13], our starting point is a generalization of the well-known Gaussian noise mechanism, which adds appropriately scaled correlated Gaussian noise to the queries. By itself, this mechanism is sufficient to guarantee privacy, but its error is too large when $n = o(|\mathcal{Q}|)$. The main insight of [NTZ13] was to use the knowledge that the database is small to reduce the error via a post-processing step. The post-processing is a form of regression: we find the vector of answers that is closest to the noisy answers while still consistent with the database size bound. (In fact the estimator is slightly more complicated and related to the hybrid estimator of Zhang [Zha13].) Intuitively, when n is small compared to the number of queries, this regression step cancels a significant fraction of the error.

¹Here and in the remainder of the introduction we ignore dependence of the error on ϵ and δ .

²We give a formal definition later.

Our first novel contribution is to analyze the error of this mechanism for arbitrary noise distributions and formulate it as a convex function of the covariance matrix of the noise. Then we write a convex program that captures the problem of finding the covariance matrix for which the performance of the mechanism is optimized on the given query workload and database size bound. We use Gaussian noise with this optimal covariance in place of the recursively constructed ad-hoc noise distribution³ from [NTZ13]. Finally, we relate the dual of the convex program to a spectral lower bound on $\text{opt}(n, \mathcal{Q})$ via the restricted invertibility principle of Bourgain and Tzafriri [BT87]. We stress that while the restricted invertibility principle was used in [NTZ13] as well, here we need a new argument which works for the optimal covariance matrix we compute and gives a smaller competitiveness ratio.

In addition to the improvement in the competitiveness ratio, our approach here is more direct and we believe it gives a better understanding of the performance of the regression-based mechanism for small databases.

2 Preliminaries

We use capital letters for matrices and lower-case letters for vectors and scalars. We use $\langle \cdot, \cdot \rangle$ for the standard inner product between vectors in \mathbb{R}^n . For a matrix $M \in \mathbb{R}^{m \times n}$ and a set $S \subseteq [n]$, we use M_S for the submatrix consisting of the columns of M indexed by elements of S . We use the notation $M \succ 0$ to denote that M is a positive definite matrix, and $M \succeq 0$ to denote that it is positive semidefinite. We use $\sigma_{\min}(M)$ for the smallest singular value of M , i.e. $\sigma_{\min}(M) \triangleq \min_x \|Mx\|_2 / \|x\|_2$. We use $\text{tr}(\cdot)$ for the trace operator, and $\|M\|_2$ for the $\ell_2 \rightarrow \ell_2$ operator norm of M , i.e. $\|M\|_2 \triangleq \max_x \|Mx\|_2 / \|x\|_2$.

The distribution of a multivariate Gaussian with mean μ and covariance Σ is denoted $N(\mu, \Sigma)$.

2.1 Histograms, the Query Matrix, and the Sensitivity Polytope

It will be convenient to encode the problem of releasing answers to linear queries using linear-algebraic notation. A common and very useful representation of a database D is the *histogram representation*: the histogram of D is a vector $x \in \mathbb{R}^{\mathcal{U}}$ such that for any $e \in \mathcal{U}$, x_e is equal to the number of copies of e in D . Notice that $\|x\|_1 = n$ and also that if x and x' are respectively the histograms of two neighboring databases D and D' , then $\|x - x'\|_1 \leq 1$ (here $\|x\|_1 = \sum_e |x_e|$ is the standard ℓ_1 norm). Linear queries are a linear transformation of x . More concretely, let us define the *query matrix* $A \in \mathbb{R}^{\mathcal{Q} \times \mathcal{U}}$ associated with a set of linear queries \mathcal{Q} by $a_{q,e} = q(e)$. Then it is easy to see that the vector Ax gives the answers to the queries \mathcal{Q} on a database D with histogram x .

Since this does not lead to any loss in generality, for the remainder of this chapter we will assume that databases are given to mechanisms as histograms, and workloads of linear queries are given as query matrices. We will identify the space of size- n databases with histograms in the scaled ℓ_1 ball $nB_1^{\mathcal{U}} \triangleq \{x \in \mathbb{R}^{\mathcal{U}} : \|x\|_1 \leq n\}$, and we will identify neighboring databases with histograms x, x' such that $\|x - x'\|_1 \leq 1$.

The *sensitivity polytope* K_A of a query matrix $A \in \mathbb{R}^{\mathcal{Q} \times \mathcal{U}}$ is the convex hull of the columns of A and the columns of $-A$. Equivalently, $K_A \triangleq AB_1^{\mathcal{U}}$, i.e. the image of the unit ℓ_1 ball in $\mathbb{R}^{\mathcal{U}}$ under multiplication by A . Notice that $nK_A = \{Ax : \|x\|_1 \leq n\}$ is the symmetric convex hull⁴ of the possible vectors of query answers to the queries in \mathcal{Q} on databases of size at most n .

2.2 Measures of Error and the Spectral Lower Bound

As our basic notion of error we will consider mean squared error. For a mechanism \mathcal{M} and a subset $X \subseteq \mathbb{R}^{\mathcal{U}}$, let us define the error with respect to the query matrix $A \in \mathbb{R}^{\mathcal{Q} \times \mathcal{U}}$ as

$$\text{err}(\mathcal{M}, X, A) \triangleq \sup_{x \in X} \left(\mathbb{E}_{\mathcal{Q}} \frac{1}{|\mathcal{Q}|} \|Ax - \mathcal{M}(A, x)\|_2^2 \right)^{1/2}.$$

³The distribution in [NTZ13] is independent of the database size bound. This could be a reason why their guarantees scale with $\log |\mathcal{Q}|$ rather than $\log n$.

⁴The symmetric convex hull of a set of points v_1, \dots, v_N is equal to the convex hull of $\pm v_1, \dots, \pm v_N$.

where the expectation is taken over the random coins of \mathcal{M} . We also write $\text{err}(\mathcal{M}, nB_1^{\mathcal{U}}, A)$ as $\text{err}(\mathcal{M}, n, A)$. The optimal error achievable by any (ε, δ) -differentially private mechanism for the query matrix A and databases of size up to n is

$$\text{opt}_{\varepsilon, \delta}(n, A) \triangleq \inf_{\mathcal{M}} \text{err}(\mathcal{M}, n, A),$$

where the infimum is taken over all (ε, δ) -differentially private mechanisms \mathcal{M} .

Arguing directly about $\text{opt}_{\varepsilon, \delta}(n, A)$ appears difficult. For this reason we use the following spectral lower bound from [NTZ13]. This lower bound was implicit in previous papers, for example [KRSU10].

Theorem 2.1 ([NTZ13]). *There exists a constant c such that for any query matrix $A \in \mathbb{R}^{\mathcal{Q} \times \mathcal{U}}$, any small enough ε , and any δ small enough with respect to ε , $\text{opt}_{\varepsilon, \delta}(n, A) \geq (c/\varepsilon) \text{SpecLB}(\varepsilon n, A)$, where*

$$\text{SpecLB}(k, A) \triangleq \max_{\substack{S \subseteq \mathcal{U} \\ |S| \leq k}} \sqrt{k/|\mathcal{Q}|} \sigma_{\min}(A_S).$$

2.3 Composition and the Gaussian Mechanism

An important basic property of differential privacy is that the privacy guarantees degrade smoothly under composition and are not affected by post-processing.

Lemma 2.1.1 ([DMNS06, DKM⁺06]). *Let $\mathcal{M}_1(\cdot)$ satisfy $(\varepsilon_1, \delta_1)$ -differential privacy, and $\mathcal{M}_2(x, \cdot)$ satisfy $(\varepsilon_2, \delta_2)$ -differential privacy for any fixed x . Then the mechanism $\mathcal{M}_2(\mathcal{M}_1(D), D)$ satisfies $(\varepsilon_1 + \varepsilon_2, \delta_1 + \delta_2)$ -differential privacy.*

A basic method to achieve (ε, δ) -differential privacy is the Gaussian mechanism. We use the following generalized variant, introduced in [NTZ13].

Theorem 2.2 ([DN03, DN04, DMNS06, NTZ13]). *Let \mathcal{Q} be a set of queries with query matrix A , and let $\Sigma \in \mathbb{R}^{\mathcal{Q} \times \mathcal{Q}}$, $\Sigma \succ 0$, be such that $a_e^T \Sigma^{-1} a_e \leq 1$ for all columns a_e of A . Then the mechanism $\mathcal{M}_{\Sigma}(A, x) = Ax + w$ where $w \sim N(0, c_{\varepsilon, \delta}^2 \Sigma)$ and $c_{\varepsilon, \delta} \triangleq \frac{0.5\sqrt{\varepsilon} + \sqrt{2\ln(1/\delta)}}{\varepsilon}$ satisfies (ε, δ) -differential privacy.*

3 The Projection Mechanism

A key element in our mechanism is the use of least squares estimation to reduce error on small databases. In this section we introduce and analyze a mechanism based on least squares estimation, similar to the hybrid estimator of [Zha13]. Essentially the same mechanism was used in [NTZ13], but the definition and analysis were tied to a particular noise distribution.

Algorithm 1 Projection Mechanism $\mathcal{M}_{\Sigma}^{\text{proj}}$

Input: (*Public*) Query matrix $A \in \mathbb{R}^{\mathcal{Q} \times \mathcal{U}}$; matrix $\Sigma \succ 0$ such that $a_e^T \Sigma^{-1} a_e \leq 1$ for all columns a_e of A .

Input: (*Private*) Histogram x of a database of size $\|x\|_1 \leq n$.

- 1: Run the generalized Gaussian mechanism (Theorem 2.2) to compute $\tilde{y} \triangleq \mathcal{M}_{\Sigma}(A, x)$;
- 2: Let Π be the orthogonal projection operator onto the span of the eigenvectors corresponding to the $\lfloor \varepsilon n \rfloor$ largest eigenvalues of Σ
- 3: Compute $\bar{y} \in n(I - \Pi)K_A$, where K_A is the sensitivity polytope of A , and \bar{y} is

$$\bar{y} = \arg \min \{ \|z - (I - \Pi)\tilde{y}\|_2^2 : z \in n(I - \Pi)K_A \}.$$

Output: Vector of answers $\Pi\tilde{y} + \bar{y}$.

As shown in [NTZ13, DNT14], Algorithm 1 can be efficiently implemented using the ellipsoid algorithm or the Frank-Wolfe algorithm.

To analyze the error of the Projection Mechanism, we use the following key lemma, which appears to be standard in statistics (we refer to [NTZ13, DNT14] for a proof). Recall that for a convex body

(compact convex set with non-empty interior) $L \subseteq \mathbb{R}^m$, the *Minkowski norm (gauge function)* is defined by $\|x\|_L \triangleq \min\{r : x \in rL\}$ for any $x \in \mathbb{R}^m$. The *polar body* is $L^\circ \triangleq \{y : \langle y, x \rangle \leq 1 \ \forall x \in L\}$ and the corresponding norm is also equal to the *support function* of L : $\|y\|_{L^\circ} \triangleq \max\{\langle y, x \rangle : x \in L\}$. When L is symmetric around 0 (i.e. $-L = L$), the Minkowski norm and support function are both norms in the usual sense.

Lemma 3.0.1 ([NTZ13, DNT14]). *Let $L \subseteq \mathbb{R}^m$ be a symmetric convex body, and let $y \in L, \tilde{y} \in \mathbb{R}^m$. Let, finally, $\bar{y} = \arg \min\{\|z - \tilde{y}\|_2^2 : z \in L\}$. We have $\|\bar{y} - y\|_2^2 \leq 4 \min\{\|\tilde{y} - y\|_2^2, \|\tilde{y} - y\|_{L^\circ}^2\}$.*

The next lemma gives our analysis of the error of the Projection Mechanism.

Lemma 3.0.2. *Assume $\Sigma \succ 0$ is such that $a_e^\top \Sigma^{-1} a_e \leq 1$ for all columns a_e of A . Then the Projection Mechanism $\mathcal{M}_\Sigma^{\text{proj}}$ in Algorithm 1 is (ε, δ) -differentially private. Moreover, for $\varepsilon = O(1)$,*

$$\text{err}(\mathcal{M}_\Sigma^{\text{proj}}, n, A) = O\left(\left(1 + \frac{\sqrt{\log |\mathcal{U}|}}{\sqrt{\log 1/\delta}}\right)^{1/2}\right) \cdot \left(\frac{c_{\varepsilon, \delta}^2}{|\mathcal{Q}|} \sum_{i \leq \varepsilon n} \sigma_i\right)^{1/2},$$

where $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{|\mathcal{Q}|}$ are the eigenvalues of Σ .

Proof. To prove the privacy guarantee, observe that the output of $\mathcal{M}_\Sigma^{\text{proj}}(A, x)$ is just a post-processing of the output of $\mathcal{M}_\Sigma(A, x)$, i.e. the algorithm does not access x except to pass it to $\mathcal{M}_\Sigma(A, x)$. The privacy then follows from Theorem 2.2 and Lemma 2.1.1.

Next we bound the error. Let $y \triangleq Ax$ be the true answers to the queries, and let $w \triangleq \tilde{y} - y \sim N(0, c_{\varepsilon, \delta}^2 \Sigma)$ be the random noise introduced by the generalized Gaussian mechanism. By the Pythagorean theorem and linearity of expectation, the expected total squared error of the projection mechanism is

$$\mathbb{E}\|\Pi\tilde{y} + \bar{y} - y\|_2^2 = \mathbb{E}\|\Pi\tilde{y} - \Pi y\|_2^2 + \mathbb{E}\|\bar{y} - (I - \Pi)y\|_2^2.$$

Above and in the remainder of the proof expectations are taken with respect to the randomness in the choice of w . We bound the two terms on the right hand side separately. We will show:

$$\mathbb{E}\|\Pi\tilde{y} - \Pi y\|_2^2 = c_{\varepsilon, \delta}^2 \sum_{i=1}^k \sigma_i, \tag{1}$$

$$\mathbb{E}\|\bar{y} - (I - \Pi)y\|_2^2 = O\left(\frac{\sqrt{\log |\mathcal{U}|}}{\sqrt{\log 1/\delta}}\right) c_{\varepsilon, \delta}^2 \sum_{i=1}^k \sigma_i. \tag{2}$$

(1) and (2) together imply the error bound in the theorem.

To prove (1), observe that $\Pi\tilde{y} - \Pi y = \Pi w \sim N(0, c_{\varepsilon, \delta}^2 \Pi \Sigma \Pi)$. By the definition of Π , the non-zero eigenvalues of $\Pi \Sigma \Pi$ are $\sigma_1, \dots, \sigma_k$ where $k \triangleq \lfloor \varepsilon n \rfloor$. We have

$$\mathbb{E}\|\Pi\tilde{y} - \Pi y\|_2^2 = c_{\varepsilon, \delta}^2 \text{tr}(\Pi \Sigma \Pi) = c_{\varepsilon, \delta}^2 \sum_{i=1}^k \sigma_i.$$

To prove (2) we appeal to Lemma 3.0.1. Define $\tilde{K} \triangleq (I - \Pi)K_A$. With $n\tilde{K}$ in the place of L , the lemma implies that

$$\mathbb{E}\|\bar{y} - (I - \Pi)y\|_2^2 \leq 4\mathbb{E}\|(I - \Pi)w\|_{(n\tilde{K})^\circ} = 4n\mathbb{E}\|(I - \Pi)w\|_{\tilde{K}^\circ}, \tag{3}$$

where we used the simple fact

$$\|(I - \Pi)w\|_{(n\tilde{K})^\circ} = \sup_{z \in n\tilde{K}} \langle (I - \Pi)w, z \rangle = n \sup_{z \in \tilde{K}} \langle (I - \Pi)w, z \rangle = n\|(I - \Pi)w\|_{\tilde{K}^\circ}.$$

\tilde{K} is the convex hull of the columns of $(I - \Pi)A$ and the columns of $-(I - \Pi)A$. For any such column $(I - \Pi)a_e$ we have

$$1 \geq a_e^\top \Sigma^{-1} a_e \geq a_e^\top (I - \Pi) \Sigma^{-1} (I - \Pi) a_e \geq \sigma_{k+1}^{-1} a_e^\top (I - \Pi) a_e.$$

The first inequality is by the assumption on Σ ; the second follows because $\Sigma^{-1} - (I - \Pi)\Sigma^{-1}(I - \Pi) \succeq 0$; the third inequality is due to the fact that the smallest eigenvalue of $(I - \Pi)\Sigma^{-1}(I - \Pi)$ restricted to the range of $I - \Pi$ is σ_{k+1}^{-1} by the choice of Π . Therefore, $\|(I - \Pi)a_e\|_2^2 \leq \sigma_{k+1} \leq \sigma_k$. Since a linear functional attains its maximum value over a polytope at a vertex, we have $\|(I - \Pi)w\|_{\tilde{K}^\circ} = \sup_{z \in \tilde{K}} \langle (I - \Pi)w, z \rangle = \max_{e \in \mathcal{U}} |\langle (I - \Pi)w, a_e \rangle|$. Each inner product $\langle (I - \Pi)w, a_e \rangle$ is a centered Gaussian random variable with variance $\mathbb{E}(\langle (I - \Pi)w, a_e \rangle)^2 = c_{\varepsilon, \delta}^2 a_e^\top (I - \Pi)\Sigma(I - \Pi)a_e$. By the choice of Π , the largest eigenvalue of $(I - \Pi)\Sigma(I - \Pi)$ is $\sigma_{k+1} \leq \sigma_k$. From this fact and the inequality $\|(I - \Pi)a_e\|_2^2 \leq \sigma_k$, we have that the variance of $\langle (I - \Pi)w, a_e \rangle$ is at most $c_{\varepsilon, \delta}^2 \sigma_k^2$. By a standard concentration argument, we can bound the expectation of the maximum absolute value of the inner products as

$$\mathbb{E}\|(I - \Pi)w\|_{\tilde{K}^\circ} = \mathbb{E}\max_{e \in \mathcal{U}} |\langle (I - \Pi)w, a_e \rangle| = O(\sqrt{\log |\mathcal{U}|})c_{\varepsilon, \delta}\sigma_k.$$

Plugging this into (3), we get

$$\mathbb{E}\|\bar{y} - (I - \Pi)y\|_2^2 = O(\sqrt{\log |\mathcal{U}|})c_{\varepsilon, \delta}n\sigma_k.$$

To show that this implies (2), observe that, by averaging, $c_{\varepsilon, \delta}n\sigma_k \leq \frac{c_{\varepsilon, \delta}n}{k} \sum_{i=1}^k \sigma_i$. Since $k = \lfloor \varepsilon n \rfloor$, $\frac{c_{\varepsilon, \delta}n}{k} = O\left(\frac{c_{\varepsilon, \delta}^2}{\sqrt{\log 1/\delta}}\right)$. This finishes the proof of (2), and, therefore, of the theorem. \square

4 Optimality of the Projection Mechanism

In this section we show that we can choose a covariance matrix Σ so that $\mathcal{M}_\Sigma^{\text{proj}}$ has nearly optimal error:

Theorem 4.1. *Let ε be a small enough constant and let $\delta = |\mathcal{U}|^{o(1)}$ be small enough with respect to ε . For any query matrix $A \in \mathbb{R}^{\mathcal{Q} \times \mathcal{U}}$, and any database size bound n , there exists a covariance matrix $\Sigma \succ 0$ such that the Projection Mechanism $\mathcal{M}_\Sigma^{\text{proj}}$ in Algorithm 1 is (ε, δ) -differentially private and has error*

$$\begin{aligned} \text{err}(\mathcal{M}, n, A) &= O((\log n)(\log 1/\delta)^{1/4}(\log |\mathcal{U}|)^{1/4}) \cdot \frac{1}{\varepsilon} \text{SpecLB}(\varepsilon n, A) \\ &= O((\log n)(\log 1/\delta)^{1/4}(\log |\mathcal{U}|)^{1/4}) \cdot \text{opt}_{\varepsilon, \delta}(n, A) \end{aligned}$$

Moreover, Σ can be computed in time polynomial in $|\mathcal{Q}|$.

Theorem 4.1 is the formal statement of Theorem 1.1. (Recall again that Algorithm 1 can be implemented in time polynomial in n , $|\mathcal{Q}|$ and $|\mathcal{U}|$, as shown in [NTZ13, DNT14].)

To prove the theorem, we optimize over the choices of Σ that ensure (ε, δ) -differential privacy, and use convex duality and the restricted invertibility principle to relate the optimal covariance to the spectral lower bound.

4.1 Minimizing the Ky Fan Norm

Recall that for an $m \times m$ matrix $\Sigma \succ 0$ with eigenvalues $\sigma_1 \geq \dots \geq \dots \geq \sigma_m$, and a positive integer $k \leq m$, the Ky Fan k -norm is defined as $\|\Sigma\|_{(k)} \triangleq \sigma_1 + \dots + \sigma_k$. The covariance matrix Σ we use in the projection mechanism will be the one achieving $\min\{\|\Sigma\|_{(k)} : a_e^\top \Sigma^{-1} a_e \leq 1 \ \forall e \in \mathcal{U}\}$, where a_e is the column of the query matrix A associated with the universe element e . This choice is directly motivated by Lemma 3.0.2. We can write this optimization problem in the following way.

$$\text{Minimize } \|X^{-1}\|_{(k)} \text{ s.t.} \tag{4}$$

$$X \succ 0 \tag{5}$$

$$\forall e \in \mathcal{U} : a_e^\top X a_e \leq 1. \tag{6}$$

The program above has a geometric meaning. For a positive definite matrix X , the set $E(X) \triangleq \{v \in \mathbb{R}^{\mathcal{Q}} : v^\top X v \leq 1\}$ is an ellipsoid centered at the origin. The constraint (6) means that $E(X)$ has to contain all columns of the query matrix A . The objective function (4) is equal to the sum of squared lengths of the k longest

major axes of $E(X)$. Therefore, we are looking for the smallest ellipsoid centered at the origin that contains the columns of A , where the “size” of the ellipsoid is the sum of squared lengths of the k longest major axes. We will not use this geometric interpretation in the rest of the paper.

We will show that (4)–(6) is a convex optimization problem. This will allow us to use general tools such as the ellipsoid method to find an optimal solution, and also to use duality theory in order to analyze the value of the optimal solution.

To show that (4)–(6) is convex we will need the following well-known result of Fan.

Lemma 4.1.1 ([Fan49]). *For any $m \times m$ real symmetric matrix Σ ,*

$$\|\Sigma\|_{(k)} = \max_{U \in \mathbb{R}^{m \times k}: U^\top U = I} \text{tr}(U^\top \Sigma U).$$

With this result in hand, we can prove that (4)–(6) is a convex optimization problem.

Lemma 4.1.2. *The objective function (4) and constraints (6) are convex over $X \succ 0$.*

Proof. The objective function and the constraints (6) are affine, and therefore convex. It remains to show that the objective (4) is also convex. Let X_1 and X_2 be two feasible solutions and define $Y = \alpha X_1 + (1 - \alpha)X_2$ for some $\alpha \in [0, 1]$. Because the matrix inverse is operator convex (see e.g. [Bha97]), $Y^{-1} \preceq \alpha X_1^{-1} + (1 - \alpha)X_2^{-1}$. Let $U \in \mathbb{R}^{m \times k}$ be such that $\text{tr}(U^\top Y^{-1} U) = \|Y^{-1}\|_{(k)}$ and $U^\top U = I$. Such a U exists by Lemma 4.1.1. We have, again using Lemma 4.1.1,

$$\begin{aligned} \|Y^{-1}\|_{(k)} &= \text{tr}(U^\top Y^{-1} U) \leq \alpha \text{tr}(U^\top X_1^{-1} U) + (1 - \alpha) \text{tr}(U^\top X_2^{-1} U) \\ &\leq \alpha \|X_1^{-1}\|_{(k)} + (1 - \alpha) \|X_2^{-1}\|_{(k)}. \end{aligned}$$

This finishes the proof. \square

Since the program (4)–(6) is convex, its optimal solution can be approximated in polynomial time within any given degree of accuracy using the ellipsoid algorithm [GLS81].

4.2 A Special Function

Before we continue, we need to introduce a somewhat complicated function of the singular values of a matrix. This function will turn out to be the objective function in a maximization problem which is dual to (4)–(6). The next lemma is needed to argue that this function is well-defined. The lemma was proved in [Nik15].

Lemma 4.1.3 ([Nik15]). *Let $\sigma_1 \geq \dots \geq \sigma_m \geq 0$ be non-negative reals, and let $k \leq m$ be a positive integer. There exists a unique integer t , $0 \leq t \leq k - 1$, such that*

$$\sigma_t > \frac{\sum_{i>t} \sigma_i}{k - t} \geq \sigma_{t+1}, \tag{7}$$

with the convention $\sigma_0 = \infty$.

We are now ready to define the function:

Definition 4.1. *Let $\Sigma \succeq 0$ be an $m \times m$ positive semidefinite matrix with singular values $\sigma_1 \geq \dots \geq \sigma_m$, and let $k \leq m$ be a positive integer. The function $h_k(\Sigma)$ is defined as*

$$h_k(\Sigma) \triangleq \sum_{i=1}^t \sigma_i^{1/2} + \sqrt{k - t} \left(\sum_{i>t} \sigma_i \right)^{1/2},$$

where t is the unique integer such that $\sigma_t > \frac{\sum_{i>t} \sigma_i}{k - t} \geq \sigma_{t+1}$.

Lemma 4.1.3 guarantees that $h_k(\Sigma)$ is a well-defined real-valued function. In the next lemma we also show that it is continuous.

Lemma 4.1.4. *The function h_k is continuous over positive semidefinite matrices with respect to the operator norm.*

Proof. Let Σ be a $m \times m$ positive semidefinite matrix with singular values $\sigma_1 \geq \dots \geq \sigma_m$ and let $t, 0 \leq t < k$, be the unique integer for which $\sigma_t > \frac{\sum_{i>t} \sigma_i}{k-t} \geq \sigma_{t+1}$. If $\frac{\sum_{i>t} \sigma_i}{k-t} > \sigma_{t+1}$, then setting δ small enough ensures that, for any Σ' such that $\|\Sigma - \Sigma'\|_2 < \delta$, $h_k(\Sigma)$ and $h_k(\Sigma')$ are computed with the same value of t . In this case, the proof of continuity follows from the continuity of the square root function. Let us therefore assume that $\frac{\sum_{i>t} \sigma_i}{k-t} = \sigma_{t+1} = \dots = \sigma_{t'} > \sigma_{t'+1}$ for some $t' \geq t+1$. Then for any integer $s \in [t, t']$,

$$\sum_{i>s} \sigma_i = \sum_{i>t} \sigma_i - (s-t)\sigma_{t+1} = (k-s)\sigma_{t+1}.$$

We then have

$$\begin{aligned} \sum_{i=1}^t \sigma_i^{1/2} + \sqrt{k-t} \left(\sum_{i>t} \sigma_i \right)^{1/2} &= \sum_{i=1}^t \sigma_i^{1/2} + (k-t)\sigma_{t+1}^{1/2} \\ &= \sum_{i=1}^s \sigma_i^{1/2} + (k-s)\sigma_{t+1}^{1/2} \\ &= \sum_{i=1}^s \sigma_i^{1/2} + \sqrt{k-s} \left(\sum_{i>s} \sigma_i \right)^{1/2}. \end{aligned} \quad (8)$$

For any Σ' such that $\|\Sigma' - \Sigma\|_2 < \delta$ for a small enough δ , we have

$$h_k(\Sigma') = \sum_{i=1}^s \sigma_i(\Sigma')^{1/2} + \sqrt{k-s} \left(\sum_{i>s} \sigma_i(\Sigma') \right)^{1/2},$$

for some integer s in $[t, t']$. Continuity then follows from (8), and the continuity of the square root function. \square

4.3 The Dual of the Ky Fan Norm Minimization Problem

Our next goal is derive a dual characterization of (4)–(6), which we will then relate to the spectral lower bound $\text{SpecLB}(k, A)$. It is useful to work with the dual, because it is a maximization problem, so to prove optimality we just need to show that any feasible solution of the dual gives a lower bound on the optimal error under differential privacy.

The next theorem gives our dual characterization in terms of the special function h_k defined in the previous section.

Theorem 4.2. *Let $A = (a_e)_{e \in \mathcal{U}} \in \mathbb{R}^{\mathcal{Q} \times \mathcal{U}}$ be a rank $|\mathcal{Q}|$ matrix, and let μ be the optimal value of (4)–(6). Then,*

$$\mu^2 = \max h_k(AQA^\top)^2 \text{ s.t.} \quad (9)$$

$$Q \succeq 0, \text{ diagonal}, \text{tr}(Q) = 1 \quad (10)$$

Since the objective of (4)–(6) is not necessarily differentiable, in order to analyze the dual and prove Theorem 4.2, we need to recall the concepts of subgradients and subdifferentials. A *subgradient* of a convex function $f: S \rightarrow \mathbb{R}$ at $x \in S$, where S is some open subset of \mathbb{R}^d , is a vector $y \in \mathbb{R}^d$ so that for every $z \in S$ we have

$$f(z) \geq f(x) + \langle z - x, y \rangle.$$

The set of subgradients of f at x is denoted $\partial f(x)$ and is known as the *subdifferential*. When f is differentiable at x , the subdifferential is a singleton set containing only the gradient $\nabla f(x)$. If f is defined by $f(x) = f_1(x) + f_2(x)$, where $f_1, f_2: S \rightarrow \mathbb{R}$, then $\partial f(x) = \partial f_1(x) + \partial f_2(x)$. A basic fact in convex analysis is

that f achieves its minimum at x if and only if $0 \in \partial f(x)$. For more information on subgradients and subdifferentials, see the classical text of Rockafellar [Roc70].

Overton and Womersley [OW93] analyzed the subgradients of functions which are a composition of a differentiable matrix-valued function with a Ky Fan norm. The special case we need also follows from the results of Lewis [Lew95].

Lemma 4.2.1 ([OW93],[Lew95]). *Let $g_k(X) \triangleq \|X^{-1}\|_{(k)}$ for a positive definite matrix $X \in \mathbb{R}^{m \times m}$. Let $\sigma_1 \geq \dots \geq \sigma_m$ be the singular values of X^{-1} and let D be the diagonal matrix with the σ_i on the diagonal. Assume that for some $r \geq k$, $\sigma_k = \dots = \sigma_r$. Then the subgradients of g_k are given by*

$$\partial g_k(X) = \text{conv}\{-U_S U_S^\top X^{-2} U_S U_S^\top : U \text{ orthonormal}, U D U^\top = X^{-1}, S \subseteq [r]\},$$

where U_S is the submatrix of U indexed by S .

We use the following well-known characterization of the convex hull of boolean vectors of Hamming weight k . For a proof, see [Sch03].

Lemma 4.2.2. *Let $V_{k,n} \triangleq \text{conv}\{v \in \{0,1\}^n : \|v\|_1 = k\}$. Then $V_{k,n} = \{v : \|v\|_1 = k, 0 \leq v_i \leq 1 \forall i\}$.*

Before we prove Theorem 4.2, we need one more technical lemma.

Lemma 4.2.3. *Let Σ be an $m \times m$ positive semidefinite matrix of rank at least k . Then there exists an $m \times m$ positive definite matrix X such that $\Sigma \in -\partial g_k(X)$, and $g_k(X) = \|X^{-1}\|_{(k)} = h_k(\Sigma)$.*

Proof. Let $r = \text{rank } \Sigma$, and let $\sigma_1 \geq \dots \geq \sigma_r$ be the non-zero singular values of Σ . Let $U D U^\top = \Sigma$ be some singular value decomposition of Σ : U is an orthonormal matrix and D is a diagonal matrix with the σ_i on the diagonal, followed by 0s.

Assume that $t, 0 \leq t < k$, is the integer (guaranteed by Lemma 4.1.3) for which $\sigma_t > \frac{\sum_{i>t} \sigma_i}{k-t} \geq \sigma_{t+1}$ and define $\alpha \triangleq \frac{\sum_{i>t} \sigma_i}{k-t}$. Since $t < k$ and we assumed Σ has rank at least k , we have $\alpha > 0$. Define

$$\sigma'_i \triangleq \begin{cases} \sigma_i & i \leq t \\ \alpha & t < i \leq r, \\ \alpha - \epsilon & i > r \end{cases}$$

and set D' be the diagonal matrix with the σ'_i on the diagonal. We set ϵ to be an arbitrary number satisfying $\alpha > \epsilon > 0$. Let us set $X \triangleq (U D' U^\top)^{-1/2}$. By Lemma 4.2.2 and the choice of t , the vector $(\sigma_{t+1}, \dots, \sigma_r)$ is an element of the polytope $\alpha V_{k-t, r-t}$. Then Σ is an element of $\text{conv}\{U_S U_S^\top X^{-2} U_S U_S^\top : S = [t] \cup T, T \subseteq \{t+1, \dots, r\}, |T| = k-t\}$. Since this set is a subset of $-\partial g_k(X)$, we have $\Sigma \in -\partial g_k(X)$. A calculation shows that $\|X^{-1}\|_{(k)} = \|(U D' U^\top)^{1/2}\|_{(k)} = \sum_{i \leq t} \sigma_i^{1/2} + (k-t)\alpha^{1/2} = h_k(\Sigma)$. This completes the proof. \square

of Theorem 4.2. We will use standard notions from the theory of convex duality. For a reference, see the book by Boyd and Vandenberghe [BV04].

Let us define $\{X : X \succ 0\}$ to be the domain for the constraints (6) and the objective function (4). This makes the constraint $X \succ 0$ implicit. The optimization problem is convex by Lemma 4.1.2. It is also always feasible: for example for r an upper bound on the Euclidean norm of the longest column of A , $\frac{1}{r}I$ is a feasible solution. Slater's condition is therefore satisfied, since the constraints are affine, and, therefore, strong duality holds.

The Lagrange dual function for (4)–(6) is

$$g(p) = \inf_{X \succ 0} \|X^{-1}\|_{(k)} + \sum_{e \in \mathcal{U}} p_e (a_e^\top X a_e - 1),$$

with dual variables $p \in \mathbb{R}^{\mathcal{U}}$, $p \geq 0$. Equivalently, we can define the diagonal matrix $P \in \mathbb{R}^{\mathcal{U} \times \mathcal{U}}$, $P \succeq 0$, with entries $p_{ee} = p_e$, and the dual function becomes

$$g(P) = \inf_{X \succ 0} \|X^{-1}\|_{(k)} + \text{tr}(A P A^\top X) - \text{tr}(P) \quad (11)$$

Since the terms $\|X^{-1}\|_{(k)}$ and $\text{tr}(APA^\top X)$ are non-negative for any $X \succ 0$, $g(P) \geq -\text{tr}(P) > -\infty$. Therefore, the effective domain $\{P : g(P) > -\infty\}$ of $g(P)$ is $\{P : P \succeq 0, \text{ diagonal}\}$. Since we have strong duality, $\mu^2 = \max\{g(P) : P \succeq 0, \text{ diagonal}\}$.

By the additivity of subgradients, a matrix X achieves the minimum in (11) if and only if $APA^\top \in -\partial g_k(X)$, where $g_k(X) = \|X^{-1}\|_{(k)}$. Consider first the case in which APA^\top has rank at least k . Then, by Lemma 4.2.3, there exists an X such that $APA^\top \in -\partial g_k(X)$ and $\|X^{-1}\|_{(k)} = h_k(APA^\top)$. Observe that, if U is an $m \times k$ matrix such that $U^\top U = I$ and $\text{tr}(U^\top X^{-1}U) = \|X^{-1}\|_{(k)}$, then

$$\text{tr}(UU^\top X^{-2}UU^\top X) = \text{tr}((U^\top X^{-2}U)(U^\top XU)) = \text{tr}(U^\top X^{-1}U) = \|X^{-1}\|_{(k)}.$$

Since, by Lemma 4.2.1 and $APA^\top \in -\partial g_k(X)$, APA^\top is a convex combination of matrices $UU^\top X^{-2}UU^\top$ with U as above, it follows that $\text{tr}(APA^\top X) = \|X^{-1}\|_{(k)}$. Then we have

$$\begin{aligned} g(P) &= \|X^{-1}\|_{(k)} + \text{tr}(APA^\top X) - \text{tr}(P) \\ &= 2\|X^{-1}\|_{(k)} - \text{tr}(P) = 2h_k(APA^\top) - \text{tr}(P). \end{aligned} \tag{12}$$

If P is such that APA^\top has rank less than k , we can reduce to the rank k case by a continuity argument. Fix any non-negative diagonal matrix P and for $\lambda \in [0, 1]$ define $P(\lambda) \triangleq \lambda P + (1 - \lambda)I$. For any $\lambda \in [0, 1]$, $AP(\lambda)A^\top$ has rank $|\mathcal{Q}|$, since AA^\top has rank $|\mathcal{Q}|$ by assumption, and, therefore, $AP(\lambda)A^\top \succeq \lambda AA^\top \succ 0$. Then, by Corollary 7.5.1. in [Roc70], and (12), we have

$$\begin{aligned} g(P) &= \lim_{\lambda \uparrow 1} g(P(\lambda)) = \lim_{\lambda \uparrow 1} [2h_k(AP(\lambda)A^\top) - \lambda \text{tr}(P) - (1 - \lambda)|\mathcal{Q}|] \\ &= 2h_k(APA^\top) - \text{tr}(P). \end{aligned}$$

The final equality follows from the continuity of h_k , proved in Lemma 4.1.4.

Let us define new variables Q and c , where $c = \text{tr}(P)$ and $Q = P/c$. Because h_k is homogeneous with exponent $1/2$, we can re-write $g(P)$ as $g(P) = g(Q, c) = 2\sqrt{c}h_k(AQA^\top) - c$. From the first-order optimality condition $\frac{\partial g}{\partial c} = 0$, we see that maximum of $g(Q, c)$ is achieved when $c = h_k(AQA^\top)^2$ and is equal to $h_k(AQA^\top)^2$. Therefore maximizing $g(P)$ over diagonal positive semidefinite P is equivalent to the optimization problem (9)–(10). Since, by strong duality, the maximum of $g(P)$ is equal to the optimal value of (4)–(6), this completes the proof. \square

4.4 Proof of Theorem 4.1

Our strategy will be to use the dual formulation in Theorem 4.2 and the restricted invertibility principle to give a lower bound on $\text{SpecLB}(k, A)$. First we state the restricted invertibility principle and a consequence of it proved in [NT15].

Theorem 4.3 ([BT87, SS10]). *Let $\epsilon \in (0, 1)$, let M be an $m \times n$ real matrix, and let W be an $n \times n$ diagonal matrix such that $W \succeq 0$ and $\text{tr}(W) = 1$. For any integer k such that $k \leq \epsilon^2 \text{tr}(MWM^\top) / \|MWM^\top\|_2$ there exists a subset $S \subseteq [n]$ of size $|S| = k$ such that $\sigma_{\min}(M_S)^2 \geq (1 - \epsilon)^2 \text{tr}(MWM^\top)$.*

For the following lemma, which is a consequence of Theorem 4.3, we need to recall the definition of the trace (nuclear) norm of a matrix M : $\|M\|_{\text{tr}}$ is equal to the sum of singular values of M .

Lemma 4.3.1 ([NT15]). *Let M be an m by n real matrix of rank r , and let $W \succeq 0$ be a diagonal matrix such that $\text{tr}(W) = 1$. Then there exists a submatrix M_S of M , $|S| \leq r$, such that $|S| \sigma_{\min}(M_S)^2 \geq c^2 \|MW^{1/2}\|_{\text{tr}}^2 / (\log r)^2$, for a universal constant $c > 0$.*

of Theorem 4.1. Given a database size n and a query matrix A , we compute the covariance matrix Σ as follows. We compute a matrix X which gives an (approximately) optimal solution to (4)–(6) for $k \triangleq \lceil \epsilon n \rceil$, and we set $\Sigma \triangleq X^{-1}$. Since (4)–(6) is a convex optimization problem, it can be solved in time polynomial in $|\mathcal{Q}|$ to any degree of accuracy using the ellipsoid algorithm [GLS81] (or the algorithm of Overton and Womersley [OW93]). By Lemma 3.0.2 and the constraints (6), $\mathcal{M}_\Sigma^{\text{proj}}$ is (ϵ, δ) -differentially private with this choice of Σ .

By Lemma 3.0.2,

$$\text{err}(\mathcal{M}_{\Sigma}^{\text{proj}}, n, A) = O\left(\left(1 + \frac{\sqrt{\log |U|}}{\sqrt{\log 1/\delta}}\right)^{1/2}\right) \cdot \frac{c_{\varepsilon, \delta}}{\sqrt{|Q|}} \|\Sigma\|_{(k)}. \quad (13)$$

By Theorem 4.2, the optimal solution Q of (9)–(10) satisfies

$$\|\Sigma\|_{(k)} = h_k(AQA^{\top}) = \sum_{i=1}^t \lambda_i^{1/2} + \sqrt{k-t} \left(\sum_{i>t} \lambda_i\right)^{1/2},$$

where $\lambda_1 \geq \dots \geq \lambda_m$ are the eigenvalues of AQA^{\top} and t , $0 \leq t < k$, is an integer such that $(k-t)\lambda_t > \sum_{i>t} \lambda_i \geq (k-t)\lambda_{t+1}$. At least one of $\sum_{i=1}^t \lambda_i^{1/2}$ and $\sqrt{k-t} \left(\sum_{i>t} \lambda_i\right)^{1/2}$ must be bounded from below by $\frac{1}{2}\|\Sigma\|_{(k)}$. Next we consider these two cases separately.

Assume first that $\sum_{i=1}^t \lambda_i^{1/2} \geq \frac{1}{2}\|\Sigma\|_{(k)}$. Let Π be the orthogonal projection operator onto the eigenspace of AQA^{\top} corresponding to $\lambda_1, \dots, \lambda_t$. Then, because $\lambda_1 \geq \dots \geq \lambda_t$ are the nonzero singular values of $\Pi A Q^{1/2}$, we have $\|\Pi A Q^{1/2}\|_{\text{tr}} = \sum_{i=1}^t \lambda_i^{1/2} \geq \frac{1}{2}\|\Sigma\|_{(k)}$. By Lemma 4.3.1 applied to the matrices $M = \Pi A$ and $W = Q$, there exists a set $S \subseteq \mathcal{U}$ of size at most $|S| \leq \text{rank } \Pi A = t < \varepsilon n$, such that

$$\begin{aligned} \text{SpecLB}(\varepsilon n, A) &\geq \sqrt{\frac{|S|}{|Q|}} \lambda_{\min}(A_S) \\ &\geq \sqrt{\frac{|S|}{|Q|}} \lambda_{\min}(\Pi A_S) \geq \frac{c \|\Pi A Q^{1/2}\|_{\text{tr}}}{(\log \varepsilon n) \sqrt{|Q|}} \geq \frac{c \|\Sigma\|_{(k)}}{2(\log \varepsilon n) \sqrt{|Q|}} \end{aligned} \quad (14)$$

for an absolute constant c .

For the second case, assume that $\sqrt{k-t} \left(\sum_{i>t} \lambda_i\right)^{1/2} \geq \frac{1}{2}\|\Sigma\|_{(k)}$. Let Π now be an orthogonal projection operator onto the eigenspace of AQA^{\top} corresponding to $\lambda_{t+1}, \dots, \lambda_m$. By the choice of t , we have

$$\frac{\text{tr}(\Pi A Q \Pi)}{\|\Pi A Q \Pi\|_2} = \frac{\sum_{i>t} \lambda_i}{\lambda_{t+1}} \geq k-t.$$

By Theorem 4.3, applied with $M = \Pi A$, $W = Q$, and $\varepsilon = \frac{1}{2}$, there exists a set $S \subseteq U$ of size $\frac{1}{4}(k-t) < k \leq \varepsilon n$ so that

$$\begin{aligned} \text{SpecLB}_2(\varepsilon n, A) &\geq \sqrt{\frac{|S|}{|Q|}} \lambda_{\min}(A_S) \\ &\geq \sqrt{\frac{|S|}{|Q|}} \lambda_{\min}(\Pi A_S) \geq \frac{\sqrt{k-t} \left(\sum_{i>t} \lambda_i\right)^{1/2}}{4\sqrt{|Q|}} \geq \frac{\|\Sigma\|_{(k)}}{8\sqrt{|Q|}}. \end{aligned} \quad (15)$$

The theorem follows from (13), the fact that at least one of (14) or (15) holds, and Theorem 2.1. \square

5 Conclusion

Several natural problems remain open. Probably the most important one is to prove results analogous to ours for worst case, rather than average, error. In that case the simple post-processing strategy of the projection mechanism will likely not be sufficient. Another interesting problem is to remove the dependence on the universe size in the competitiveness ratio. It is plausible that this can be done with the projection mechanism and a well-chosen Gaussian noise distribution, but we would need tighter lower bounds, possibly based on fingerprinting codes as in [BUV13].

Acknowledgments

The author would like to thank the anonymous reviewers of ICALP 2015 for helpful comments.

References

- [Bha97] Rajendra Bhatia. *Matrix analysis*, volume 169 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1997.
- [BLR08] Avrim Blum, Katrina Ligett, and Aaron Roth. A learning theory approach to non-interactive database privacy. In *STOC '08: Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 609–618, New York, NY, USA, 2008. ACM.
- [BT87] J. Bourgain and L. Tzafriri. Invertibility of large submatrices with applications to the geometry of banach spaces and harmonic analysis. *Israel journal of mathematics*, 57(2):137–224, 1987.
- [BUV13] Mark Bun, Jonathan Ullman, and Salil Vadhan. Fingerprinting codes and the price of approximate differential privacy. *arXiv preprint arXiv:1311.3158*, 2013.
- [BV04] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, Cambridge, 2004.
- [CSS10] T-H. Hubert Chan, Elaine Shi, and Dawn Song. Private and continual release of statistics. In *ICALP*, 2010.
- [DKM⁺06] Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. 4004:486–503, 2006.
- [DMNS06] C. Dwork, F. Mcsherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *TCC*, 2006.
- [DMT07] Cynthia Dwork, Frank McSherry, and Kunal Talwar. The price of privacy and the limits of lp decoding. In *STOC*, pages 85–94, 2007.
- [DN03] Irit Dinur and Kobbi Nissim. Revealing information while preserving privacy. pages 202–210, 2003.
- [DN04] Cynthia Dwork and Kobbi Nissim. Privacy-preserving datamining on vertically partitioned databases. In *CRYPTO*, pages 528–544, 2004.
- [DNPR10] Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N. Rothblum. Differential privacy under continual observation. In Leonard J. Schulman, editor, *STOC*, pages 715–724. ACM, 2010.
- [DNR⁺09] Cynthia Dwork, Moni Naor, Omer Reingold, Guy N Rothblum, and Salil Vadhan. On the complexity of differentially private data release: efficient algorithms and hardness results. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, pages 381–390. ACM, 2009.
- [DNT14] Cynthia Dwork, Aleksandar Nikolov, and Kunal Talwar. Using convex relaxations for efficiently and privately releasing marginals. In Siu-Wing Cheng and Olivier Devillers, editors, *30th Annual Symposium on Computational Geometry, SOCG'14, Kyoto, Japan, June 08 - 11, 2014*, page 261. ACM, 2014.
- [DRV10] Cynthia Dwork, Guy N. Rothblum, and Salil Vadhan. Boosting and differential privacy. In *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science, FOCS '10*, pages 51–60, Washington, DC, USA, 2010. IEEE Computer Society.
- [Fan49] Ky Fan. On a theorem of Weyl concerning eigenvalues of linear transformations. I. *Proc. Nat. Acad. Sci. U. S. A.*, 35:652–655, 1949.
- [GHRU11] Anupam Gupta, Moritz Hardt, Aaron Roth, and Jonathan Ullman. Privately releasing conjunctions and the statistical query barrier. In *STOC*, pages 803–812, 2011.

- [GKS08] Srivatsava Ranjit Ganta, Shiva Prasad Kasiviswanathan, and Adam Smith. Composition attacks and auxiliary information in data privacy. In Ying Li, Bing Liu, and Sunita Sarawagi, editors, *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Las Vegas, Nevada, USA, August 24-27, 2008*, pages 265–273. ACM, 2008.
- [GLS81] M. Grötschel, L. Lovász, and A. Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, 1981.
- [GRU12] Anupam Gupta, Aaron Roth, and Jonathan Ullman. Iterative constructions and private data release. In *TCC*, pages 339–356, 2012.
- [HLM12] Moritz Hardt, Katrina Ligett, and Frank McSherry. A simple and practical algorithm for differentially private data release. In *NIPS*, 2012. To appear.
- [HR10] M. Hardt and G. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. *Proc. 51st Foundations of Computer Science (FOCS). IEEE*, 2010.
- [KRSU10] S.P. Kasiviswanathan, M. Rudelson, A. Smith, and J. Ullman. The price of privately releasing contingency tables and the spectra of random matrices with correlated rows. In *Proceedings of the 42nd ACM symposium on Theory of computing*, pages 775–784. ACM, 2010.
- [Lew95] A. S. Lewis. The convex analysis of unitarily invariant matrix functions. *J. Convex Anal.*, 2(1-2):173–183, 1995.
- [MN12] S. Muthukrishnan and Aleksandar Nikolov. Optimal private halfspace counting via discrepancy. In Howard J. Karloff and Toniann Pitassi, editors, *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 1285–1292. ACM, 2012.
- [Nik15] Aleksandar Nikolov. Randomized rounding for the largest j -simplex problem. *To Appear in STOC 15.*, 2015.
- [NT15] Aleksandar Nikolov and Kunal Talwar. Approximating hereditary discrepancy via small width ellipsoids. In Piotr Indyk, editor, *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 324–336. SIAM, 2015.
- [NTZ13] Aleksandar Nikolov, Kunal Talwar, and Li Zhang. The geometry of differential privacy: the sparse and approximate cases. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013*, pages 351–360. ACM, 2013.
- [OW93] M. L. Overton and R. S. Womersley. Optimality conditions and duality theory for minimizing sums of the largest eigenvalues of symmetric matrices. *Math. Programming*, 62(2, Ser. B):321–357, 1993.
- [Roc70] R. Tyrrell Rockafellar. *Convex analysis*. Princeton Mathematical Series, No. 28. Princeton University Press, Princeton, N.J., 1970.
- [RR10] Aaron Roth and Tim Roughgarden. Interactive privacy via the median mechanism. In *Proceedings of the 42nd ACM symposium on Theory of computing*, STOC ’10, pages 765–774, New York, NY, USA, 2010. ACM.
- [Sch03] Alexander Schrijver. *Combinatorial optimization. Polyhedra and efficiency. Vol. B*, volume 24 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2003. Matroids, trees, stable sets, Chapters 39–69.
- [SS10] D.A. Spielman and N. Srivastava. An elementary proof of the restricted invertibility theorem. *Israel Journal of Mathematics*, pages 1–9, 2010.

- [Ull13] Jonathan Ullman. Answering $n^{2+o(1)}$ counting queries with differential privacy is hard. In *STOC*, 2013.
- [XWG10] Xiaokui Xiao, Guozhang Wang, and Johannes Gehrke. Differential privacy via wavelet transforms. In *ICDE*, pages 225–236, 2010.
- [Zha13] Li Zhang. Nearly optimal minimax estimator for high dimensional sparse linear regression. *Annals of Statistics*, 2013. To appear.